Australian Bureau of Statistics

1346.0.55.002 - Time Series Analysis Frequently Asked Questions, 2003

Latest ISSUE Released at 11:30 AM (CANBERRA TIME) 03/02/2006

Summary

Main Features

This document was added or updated on 28/03/2017.

- **The Basics**
  An introduction to the principles of time series analysis and seasonal adjustment
- **The Process of Seasonal Adjustment**
  A look at how economic time series are seasonally adjusted
- **Issues with Seasonal Adjustment**
  Some of the problems encountered during seasonal adjustment
- **Seasonal Adjustment Methods**
  Methods used by the ABS and other statistical agencies to analyse time series
- **Further Reading**
  More reading material on time series analysis

About This Release

The Time series Analysis Frequently Asked Questions is a web page with links to several other pages that contain common questions and answers related to time series and the seasonal adjustment process. The topic headings for each of the different pages are: The basics; The process of seasonal adjustment; Issues with seasonal adjustment; Seasonal adjustment methods; and Further reading.

Time Series Analysis: The Basics

WHAT IS A TIME SERIES?

A time series is a collection of observations of well-defined data items obtained through repeated measurements over time. For example, measuring the value of retail sales each month of the year would comprise a time series. This is because sales revenue is well defined, and consistently measured at equally spaced intervals. Data collected irregularly or only once are not time series.

An observed time series can be decomposed into three components: the trend (long term direction), the seasonal (systematic, calendar related movements) and the irregular (unsystematic, short term fluctuations).

WHAT ARE STOCK AND FLOW SERIES?

Time series can be classified into two different types: stock and flow.

A stock series is a measure of certain attributes at a point in time and can be thought of as “stocktakes”. For example, the **Monthly Labour Force Survey** is a stock measure because it takes stock of whether a person was employed in the reference week.

Flow series are series which are a measure of activity over a given period. For example, surveys of **Retail Trade** activity. Manufacturing is also a flow measure because a certain amount is produced each day, and then these amounts are summed to give a total value for production for a given reporting period.

The main difference between a stock and a flow series is that flow series can contain effects related to the calendar (trading day effects). Both types of series can still be seasonally adjusted using the same seasonal adjustment process.

WHAT ARE SEASONAL EFFECTS?

A seasonal effect is a systematic and calendar related effect. Some examples include the sharp escalation in most Retail series which occurs around December in response to the Christmas period, or an increase in water consumption in summer due to warmer weather. Other seasonal effects include trading day effects (the number of working or trading days in a given month differs from year to year which will impact upon the level of activity in that month) and moving holiday (the timing of holidays such as Easter varies, so the effects of the holiday will be experienced in different periods each year).

WHAT IS SEASONAL ADJUSTMENT AND WHY DO WE NEED IT?

Seasonal adjustment is the process of estimating and then removing from a time series influences that are systematic and calendar related. Observed data needs to be seasonally adjusted as seasonal effects can conceal both the true underlying movement in the series, as well as certain non-seasonal characteristics which may be of interest to analysts.

WHY CAN'T WE JUST COMPARE ORIGINAL DATA FROM THE SAME PERIOD IN

EACH YEAR?

A comparison of original data from the same period in each year does not completely remove all seasonal effects. Certain holidays such as Easter and Chinese New Year fall in different periods in each year, hence they will distort observations. Also, year to year values will be biased by any changes in seasonal patterns that occur over time. For example, consider a comparison between two consecutive March months i.e. compare the level of the original series observed in March for 2000 and 2001. This comparison ignores the moving holiday effect of Easter. Easter occurs in April for most years but if Easter falls in March, the level of activity can vary greatly for that month for some series. This distorts the original estimates. A comparison of these two months will not reflect the underlying pattern of the data. The comparison also ignores trading day effects. If the two consecutive months of March have different composition of trading days, it might reflect different levels of activity in original terms even though the underlying level of activity is unchanged. In a similar way, any changes to seasonal patterns might also be ignored. The original estimates also contains the influence of the irregular component. If the magnitude of the irregular component of a series is strong compared with the magnitude of the trend component, the underlying direction of the series can be distorted.

However, the major disadvantage of comparing year to year original data, is lack of precision and time delays in the identification of turning points in a series. Turning points occur when the direction of underlying level of the series changes, for example when a consistently decreasing series begins to rise steadily. If we compare year apart data in the original series, we may miss turning points occurring during the year. For example, if March 2001 has a higher original estimate than March 2000, by comparing these year apart values, we might conclude that the level of activity has increased during the year. However, the series might have increased up to September 2000 and then started to decrease steadily.

WHEN IS SEASONAL ADJUSTMENT INAPPROPRIATE?

When a time series is dominated by the trend or irregular components, it is nearly impossible to identify and remove what little seasonality is present. Hence seasonally adjusting a non-seasonal series is impractical and will often introduce an artificial seasonal element.

WHAT IS SEASONALITY?

The seasonal component consists of effects that are reasonably stable with respect to timing, direction and magnitude. It arises from systematic, calendar related influences such as:

- **Natural Conditions**
  weather fluctuations that are representative of the season (uncharacteristic weather patterns such as snow in summer would be considered irregular influences)
- **Business and Administrative procedures**
  start and end of the school term
- **Social and Cultural behaviour**
  Christmas

It also includes calendar related systematic effects that are not stable in their annual timing or are caused by variations in the calendar from year to year, such as:

- **Trading Day Effects**

the number of occurrences of each of the day of the week in a given month will differ from year to year
- There were 4 weekends in March in 2000, but 5 weekends in March of 2002

- **<u>Moving Holiday Effects</u>**
  holidays which occur each year, but whose exact timing shifts
  - Easter, Chinese New Year

HOW DO WE IDENTIFY SEASONALITY?

Seasonality in a time series can be identified by regularly spaced peaks and troughs which have a consistent direction and approximately the same magnitude every year, relative to the trend. The following diagram depicts a strongly seasonal series. There is an obvious large seasonal increase in December retail sales in New South Wales due to Christmas shopping. In this example, the magnitude of the seasonal component increases over time, as does the trend.

Figure 1: Monthly Retail Sales in New South Wales (NSW) Retail Department Stores

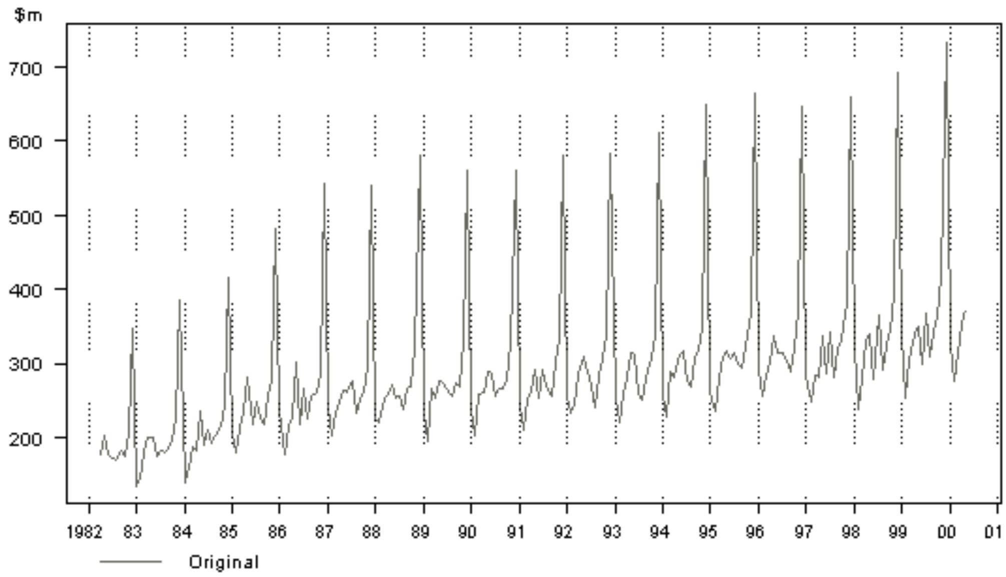

WHAT IS AN IRREGULAR?

The irregular component (sometimes also known as the residual) is what remains after the seasonal and trend components of a time series have been estimated and removed. It results from short term fluctuations in the series which are neither systematic nor predictable. In a highly irregular series, these fluctuations can dominate movements, which will mask the trend and seasonality. The following graph is of a highly irregular time series:

Figure 2: Monthly Value of Building Approvals, Australian Capital Territory (ACT)

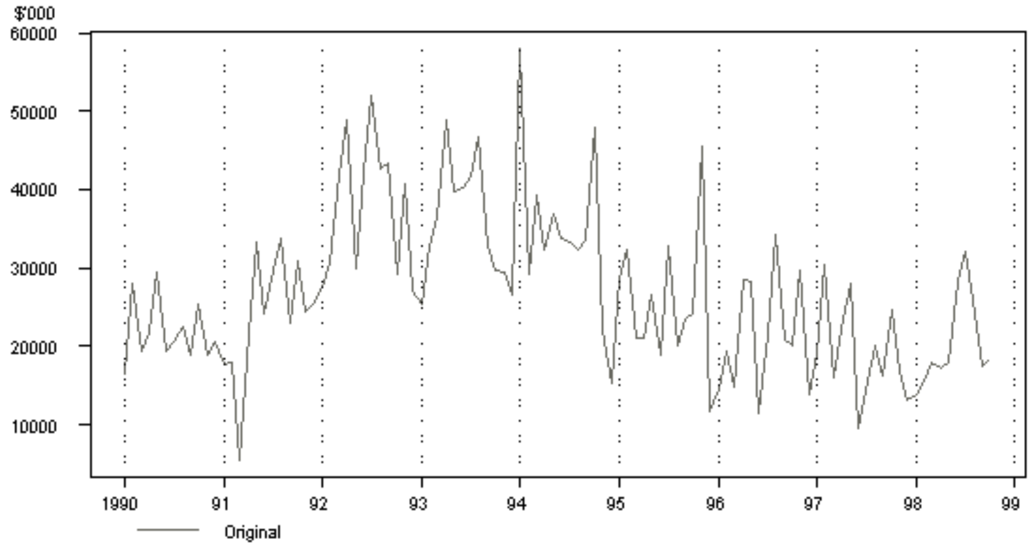

WHAT IS THE TREND?

The ABS trend is defined as the 'long term' movement in a time series without calendar related and irregular effects, and is a reflection of the underlying level. It is the result of influences such as population growth, price inflation and general economic changes. The following graph depicts a series in which there is an obvious upward trend over time:

Figure 3: Quarterly Gross Domestic Product

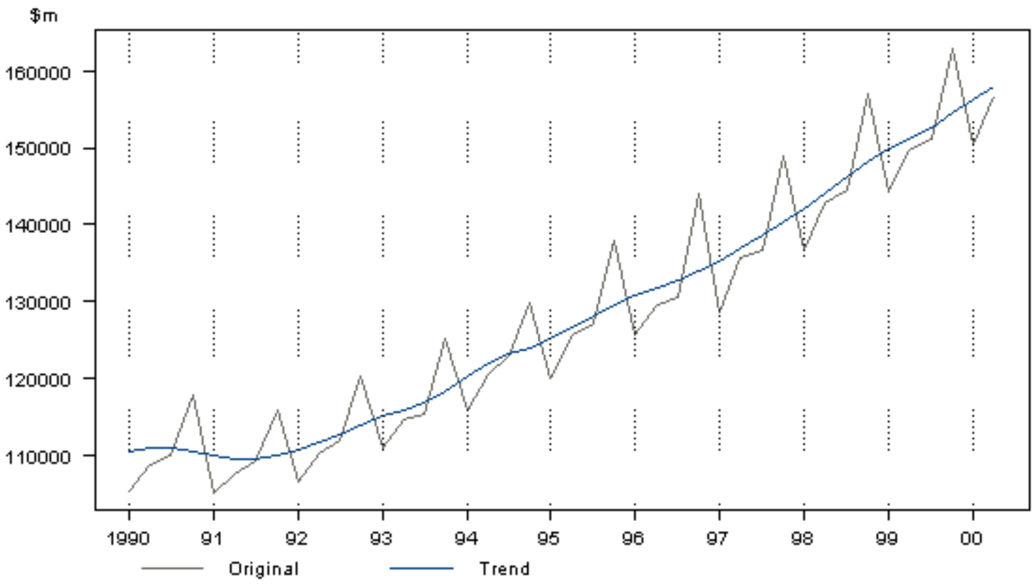

WHAT ARE THE UNDERLYING MODELS USED TO DECOMPOSE THE OBSERVED TIME SERIES?

Decomposition models are typically additive or multiplicative, but can also take other forms such as pseudo-additive.

Additive Decomposition

In some time series, the amplitude of both the seasonal and irregular variations do not change as the level of the trend rises or falls. In such cases, an additive model is

appropriate.

In the additive model, the observed time series (O_t) is considered to be the sum of three independent components: the seasonal S_t, the trend T_t and the irregular I_t.

$$\text{Observed series} = \text{Trend} + \text{Seasonal} + \text{Irregular}$$

That is

$$O_t = T_t + S_t + I_t$$

Each of the three components has the same units as the original series. The seasonally adjusted series is obtained by estimating and removing the seasonal effects from the original time series. The estimated seasonal component is denoted by $\hat{S}_t$ The seasonally adjusted estimates can be expressed by:

$$\begin{aligned}\text{Seasonally adjusted series} &= \text{Observed series} - \text{Seasonal} \\ &= \text{Trend} + \text{Irregular}\end{aligned}$$

In symbols,

$$\begin{aligned}SA_t &= O_t - \hat{S}_t \\ &= T_t + I_t\end{aligned}$$

The following figure depicts a typically additive series. The underlying level of the series fluctuates but the magnitude of the seasonal spikes remains approximately stable.

Figure 4: General Government and Other Current Transfers to Other Sectors

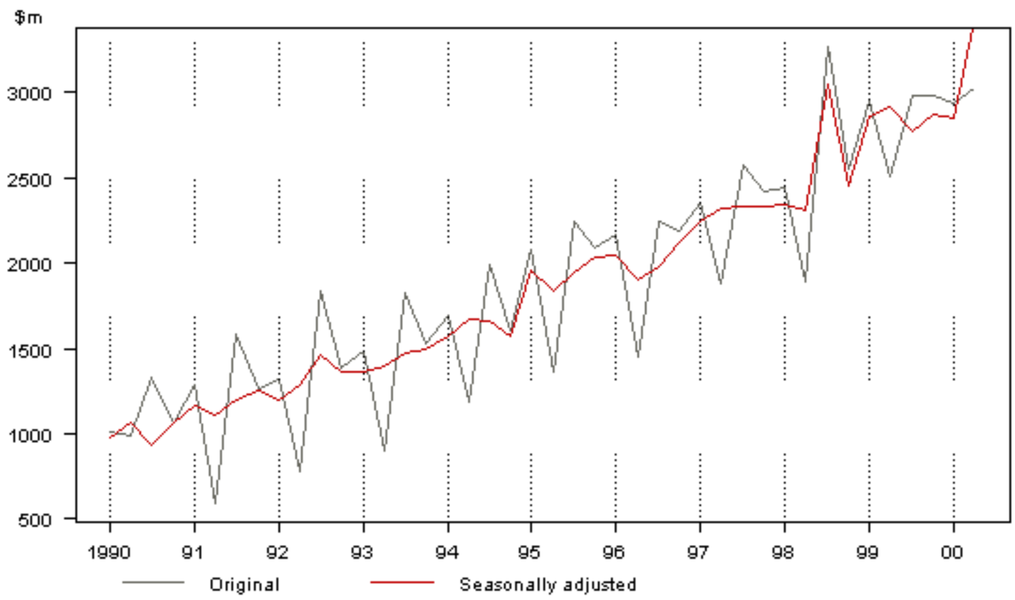

Multiplicative Decomposition

In many time series, the amplitude of both the seasonal and irregular variations increase as the level of the trend rises. In this situation, a multiplicative model is usually appropriate.

In the multiplicative model, the original time series is expressed as the product of trend,

seasonal and irregular components.

$$\text{Observed series} = \text{Trend} \times \text{Seasonal} \times \text{Irregular}$$

or

$$O_t = T_t \times S_t \times I_t$$

The seasonally adjusted data then becomes:

$$\begin{aligned}\text{Seasonally Adjusted series} &= \text{Observed} \div \text{Seasonal} \\ &= \text{Trend} \times \text{Irregular}\end{aligned}$$

or

$$\begin{aligned}SA_t &= O_t / S_t \\ &= T_t \times I_t\end{aligned}$$

Under this model, the trend has the same units as the original series, but the seasonal and irregular components are unitless factors, distributed around 1.

Most of the series analysed by the ABS show characteristics of a multiplicative model. As the underlying level of the series changes, the magnitude of the seasonal fluctuations varies as well.

Figure 5: Monthly NSW ANZ Job Advertisements

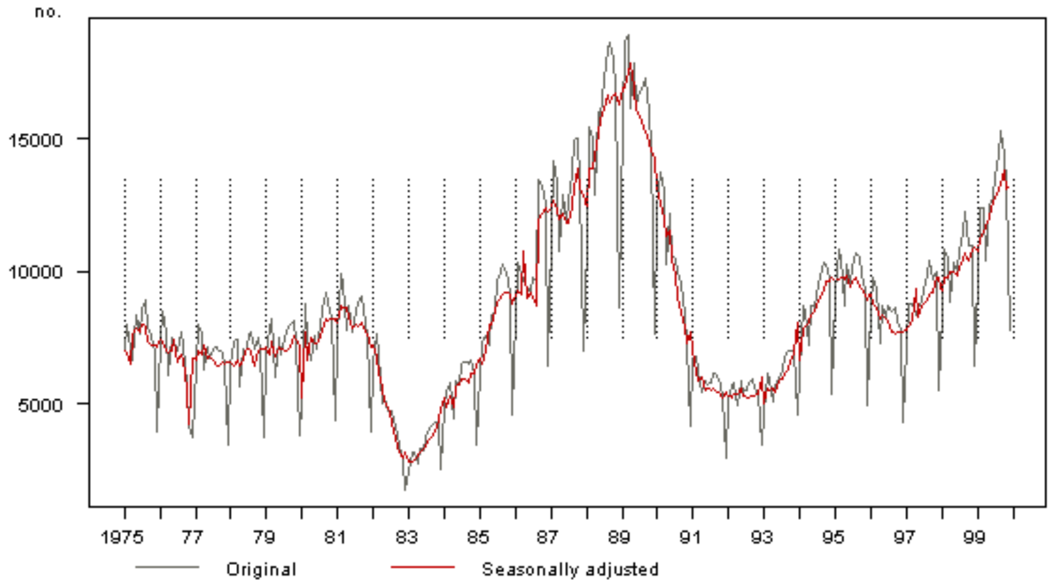

Pseudo-Additive Decomposition

The multiplicative model cannot be used when the original time series contains very small or zero values. This is because it is not possible to divide a number by zero. In these cases, a pseudo additive model combining the elements of both the additive and multiplicative models is used. This model assumes that seasonal and irregular variations are both dependent on the level of the trend but independent of each other.

The original data can be expressed in the following form:

$$
\begin{aligned}
O_t &= T_t + T_t \times (S_t - 1) + T_t \times (I_t - 1) \\
&= T_t \times (S_t + I_t - 1)
\end{aligned}
$$

The pseudo-additive model continues the convention of the multiplicative model to have both the seasonal factor S_t and the irregular factor I_t centred around one. Therefore we need to subtract one from S_t and I_t to ensure that the terms T_t x (S_t - 1) and T_t x (I_t - 1) are centred around zero. These terms can be interpreted as the additive seasonal and additive irregular components respectively and because they are centred around zero the original data O_t will be centred around the trend values T_t .

The seasonally adjusted estimate is defined to be:

$$
\begin{aligned}
SA_t &= O_t - \hat{T}_t \times (\hat{S}_t - 1) \\
&= T_t \times I_t
\end{aligned}
$$

where $\hat{T}_t$ and $\hat{S}_t$ are the trend and seasonal component estimates. In the pseudo-additive model, the trend has the same units as the original series, but the seasonal and irregular components are unitless factors, distributed around 1.

An example of series that requires a pseudo-additive decomposition model is shown below. This model is used as cereal crops are only produced during certain months, with crop production being virtually zero for one quarter each year.

Figure 6: Quarterly Gross Value for the Production of Cereal Crops

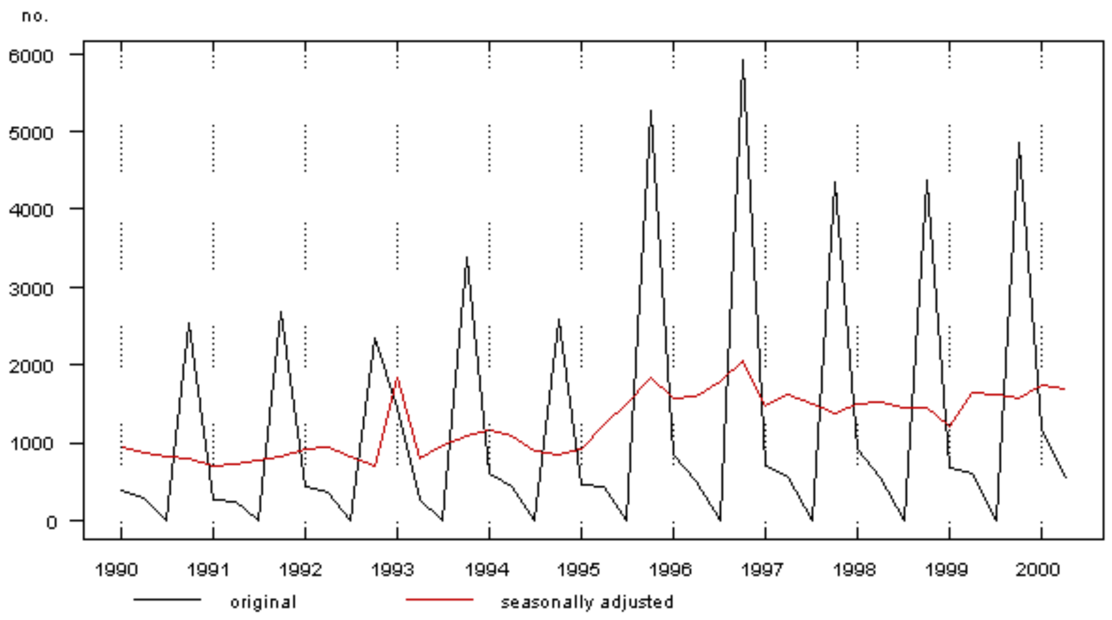

Example: Shiskin Decomposition

The Shiskin decomposition gives graphs of the original series, seasonally adjusted series, trend series, residual (irregular) factors and the between month (seasonal) and within month (trading day) factors that are combined to form the combined adjustment factors. The residual (irregular) factors are found by dividing the seasonally adjusted series by the trend series. Figure 7 shows a Shiskin decomposition for the Australian Retail series.

Figure 7: Shiskin decomposition for Australian Total Retail Turnover, May 1990 to May 2000

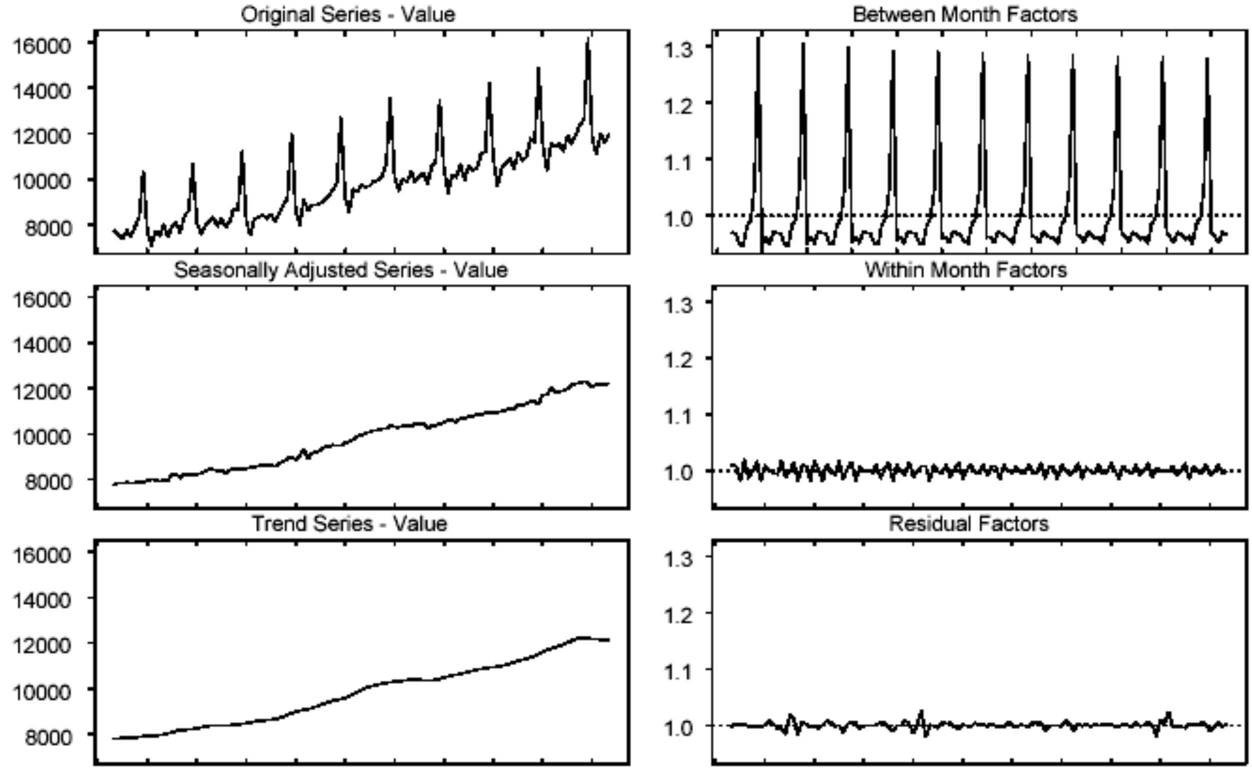

HOW DO I KNOW WHICH DECOMPOSITION MODEL TO USE?

To choose an appropriate decomposition model, the time series analyst will examine a graph of the original series and try a range of models, selecting the one which yields the most stable seasonal component. If the magnitude of the seasonal component is relatively constant regardless of changes in the trend, an additive model is suitable. If it varies with changes in the trend, a multiplicative model is the most likely candidate. However if the series contains values close or equal to zero, and the magnitude of seasonal component appears to be dependent upon the trend level, then pseudo-additive model is most appropriate.

WHAT IS A SEASONAL AND IRREGULAR (SI) CHART?

Once the trend component is estimated, it can be removed from the original data, leaving behind the combined seasonal and irregular components or SIs. A seasonal and irregular or SI chart graphically presents the SI's for particular months or quarters in the series span.

The following graph is an SI chart for a monthly series, using a multiplicative decomposition model.

Figure 8: Seasonal and Irregular (SI) Chart - Value of Building Approvals, ACT

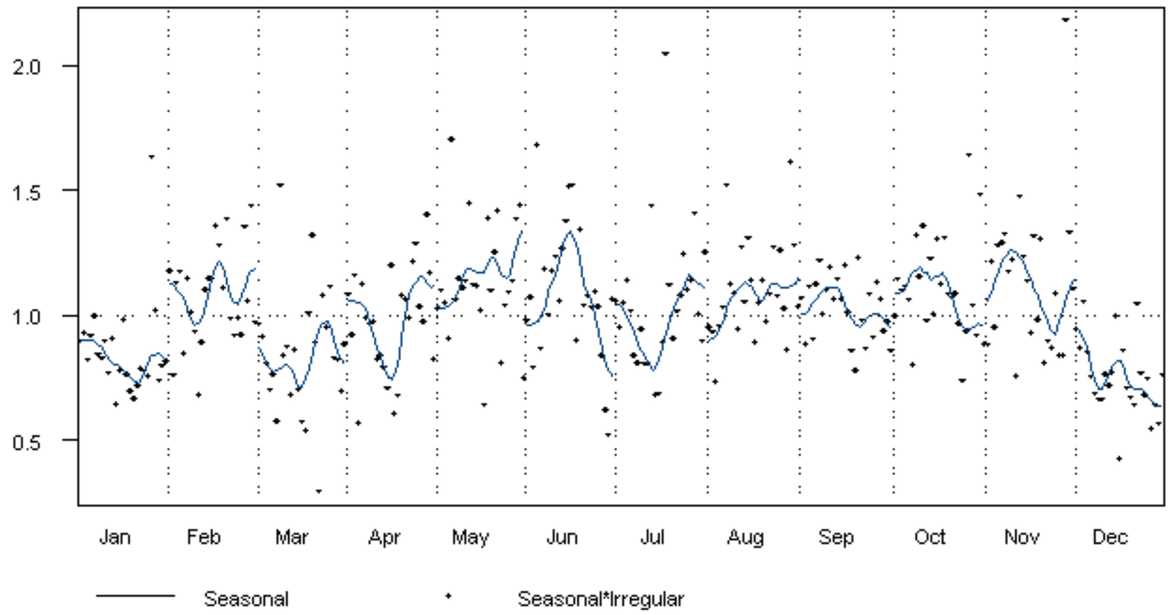

The points represent the SIs obtained from the time series, while the solid line shows the seasonal component. The seasonal component is calculated by smoothing the SI's, to remove irregular influences.

SI charts are useful in determining whether short-term movements are caused by seasonal or irregular influences. In the graph above, the SIs can be seen to fluctuate erratically, which indicates the time series under analysis is dominated by its irregular component.

SI charts are also used to identify seasonal breaks, moving holiday patterns and extreme values in a time series.

This page first published 14 November 2005, last updated 1 March 2023

Time Series Analysis: The Process of Seasonal Adjustment

WHAT ARE THE TWO MAIN PHILOSOPHIES OF SEASONAL ADJUSTMENT?

The two main philosophies for seasonal adjustment are the model based method and the filter based method.

Filter based methods

This method applies a set of fixed filters (moving averages) to decompose the time series into a trend, seasonal and irregular component.

The underlying notion is that economic data is made up of a range of cycles, including business cycles (the trend), seasonal cycles (seasonality) and noise (the irregular component). A filter essentially removes or reduces the strength of certain cycles from the input data.

To produce a seasonally adjusted series from data collected monthly, events that occur every 12, 6, 4, 3, 2.4 and 2 months need to be removed. These correspond to seasonal frequencies of 1, 2, 3, 4, 5 and 6 cycles per year. The longer non-seasonal cycles are considered to be part of the trend and the shorter non-seasonal cycles form the irregular. However the boundary between the trend and irregular cycles can vary with the length of the filter used to obtain the trend. In ABS seasonal adjustment, cycles which contribute significantly to the trend are typically larger than about 8 months for monthly series and 4 quarters for quarterly series.

The trend, seasonal and irregular components do not need explicit individual models. The irregular component is defined as what remains after the trend and seasonal components have been removed by filters. Irregulars do not display white noise characteristics.

Filter based methods are often known as X11 style methods. These include X11 (developed by U.S Census Bureau), X11ARIMA (developed by Statistics Canada), X12ARIMA (developed by U.S Census Bureau), STL, SABL and SEASABS (the package used by the ABS)..

Computational differences between various methods in X11 family are chiefly the result of different techniques used at the ends of the time series. For example, some methods use asymmetric filters at the ends, while other methods extrapolate the time series and apply symmetric filters to the extended series.

Model based methods

This approach requires the trend, seasonal and irregular components of the time series to be modelled separately. It assumes the irregular component is “white noise” - that is all cycle lengths are equally represented. The irregulars have zero mean and a constant variance. The seasonal component has its own noise element.

Two widely used software packages which apply model based methods are STAMP and SEATS/TRAMO (developed by the Bank of Spain).

Major computational differences between the various model based methods are usually due to model specifications. In some cases, the components are modelled directly. Other methods require the original time series to be modelled first, and the component models decomposed from that.

WHAT IS A FILTER?

Filters can be used to decompose a time series into a trend, seasonal and irregular component. Moving averages are a type of filter that successively average a shifting time span of data in order to produce a smoothed estimate of a time series. This smoothed series can be considered to have been derived by running an input series through a process which filters out certain cycles. Consequently, a moving average is often referred to as a filter.

The basic process involves defining a set of weights of length $m_1 + m_2 + 1$ as:

$$w_{-m_1}, w_{-(m_1-1)}, \ldots, w_{-1}, w_0, w_1, \ldots, w_{m_2-1}, w_{m_2}$$

Note: a symmetric set of weights has $m_1=m_2$ and $w_j = w_{-j}$

A filtered value at time t can be calculated by

$$Z_t = \sum_{j=-m_1}^{m_2} w_j Y_{t+j}$$

where Y_t describes the value of the time series at time t.

For example, consider the following series:

10	12	8	10	12	14	6	10

Using a simple 3 term symmetric filter (i.e. $m_1= m_2= 1$ and all weights are 1/3), the first term of the smoothed series is obtained by applying the weights to the first three terms of the original series:

$$\tfrac{1}{3} \times 10 + \tfrac{1}{3} \times 12 + \tfrac{1}{3} \times 8 = 10$$

The second smoothed value is produced by applying the weights to the second, third and fourth terms in the original series:

$$\tfrac{1}{3} \times 12 + \tfrac{1}{3} \times 8 + \tfrac{1}{3} \times 10 = 10$$

And so on…

10	12	8	10	12	14	6	10

$$\left(\tfrac{1}{3} \quad \tfrac{1}{3} \quad \tfrac{1}{3}\right) \rightarrow$$

The smoothed series becomes

10	10	10	12	10.3	10

WHAT IS THE END POINT PROBLEM?

Reconsider the series:

10	12	8	10	12	14	6	10

This series contains 8 terms. However, the smoothed series obtained by applying symmetric filter to the original data contains only 6 terms:

10	10	10	12	10.3	10

This is because there is insufficient data at the ends of the series to apply a symmetric filter. The first term of the smoothed series is a weighted average of three terms, centered on the second term of the original series. A weighted average centered on the first term of the original series cannot be obtained as data before this point is not available. Similarly, it is not possible to calculate a weighted average centered on the last term of the series, as there is no data after this point.

For this reason, symmetric filters cannot be used at either end of a series. This is known as the end point problem. Time series analysts can use asymmetric filters to produce smoothed estimates in these regions. In this case, the smoothed value is calculated ‘off centre’, with the average being determined using more data from one side of the point than the other according to what is available. Alternatively, modelling techniques may be used to extrapolate the time series and then apply symmetric filters to the extended series.

HOW DO WE DECIDE WHICH FILTER TO USE?

The time series analyst chooses an appropriate filter based on its properties, such as which cycles the filter removes when applied. The properties of a filter can be investigated using a gain function.

WHAT IS A GAIN FUNCTION?

Gain functions are used to examine the effect of a filter at a given frequency on the amplitude of a cycle for a particular time series. The following diagram is the gain function for the symmetric 3 term filter we studied earlier:

Figure 1: Gain Function for Symmetric 3 Term Filter

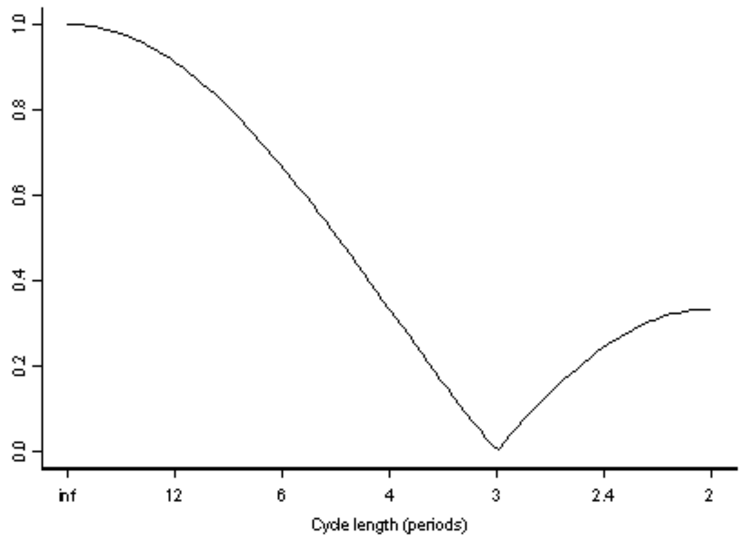

The horizontal axis represents the length of an input cycle relative to the period between observation points in the original time series. So an input cycle of length 2 is completed in 2 periods, which represents 2 months for a monthly series, and 2 quarters for a quarterly series. The vertical axis shows the amplitude of the output cycle relative to an input cycle.

This filter reduces the strength of 3 period cycles to zero. That is, it completely removes cycles of approximately this length. This means that for a time series where data is collected monthly, any seasonal effects which occur quarterly will be eliminated by applying this filter to the original series.

WHAT IS A PHASE SHIFT?

A phase shift is the time shift between the filtered cycle and the unfiltered cycle. A positive phase shift means that the filtered cycle is shifted backwards and a negative phase shift it is shifted forwards in time.

Phase shifting occurs when timing of turning points is distorted, for example when the moving average is placed off-centre by the asymmetric filters. That is they will occur either earlier or later in the filtered series, than in the original. Odd length symmetric moving averages (as used by the ABS), where the result is centrally placed, do not cause time phase shifting. It is important for filters used to derive the trend to retain the time phase, and hence the timing of any turning points.

Figures 2 and 3 show the effects of applying a 2x12 symmetric moving average which is off-centre. The continuous curves represent the initial cycles and the broken curves represents the output cycles after applying the moving average filter.

Figure 2: 24 Month Cycle, Phase = -5.5 months Amplitude = 63%

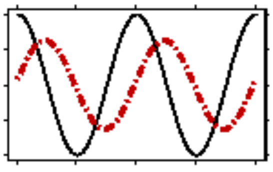

Figure 3: 8 Month Cycle, Phase = -1.5 months Amplitude = 22%

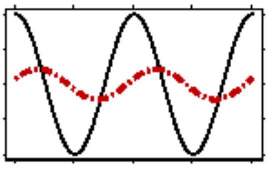

WHAT ARE HENDERSON MOVING AVERAGES?

Henderson moving averages are filters which were derived by Robert Henderson in 1916 for use in actuarial applications. They are trend filters, commonly used in time series analysis to smooth seasonally adjusted estimates in order to generate a trend estimate. They are used in preference to simpler moving averages because they can reproduce polynomials of up to degree 3, thereby capturing trend turning points.

The ABS uses Henderson moving averages to produce trend estimates from a seasonally adjusted series. The trend estimates published by the ABS are typically derived using a 13 term Henderson filter for monthly series, and a 7 term Henderson filter for quarterly series.

Henderson filters can be either symmetric or asymmetric. Symmetric moving averages can be applied at points which are sufficiently far away from the ends of a time series. In this case, the smoothed value for a given point in the time series is calculated from an equal number of values on either side of the data point.

To obtain the weights, a compromise is struck between the two characteristics generally expected of a trend series. These are that the trend should be able to represent a wide range of curvatures and that it should also be as smooth as possible.

The weighting patterns for a range of symmetric Henderson moving averages are given in the following table:

Filter Length	Symmetric Weighting Pattern for Henderson Moving Average
5 Term	(-0.073, 0.294, 0.558, 0.294, -0.073)
7 Term	(-0.059, 0.059, 0.294, 0.412, 0.294, 0.059, -0.059)
9 Term	(-0.041, -0.010, 0.119, 0.267, 0.330, 0.267, 0.119, -0.010, -0.041)
13 Term	(-0.019, -0.028, 0.0, 0.066, 0.147, 0.214, 0.240, 0.214, 0.147, 0.066, 0.0, -0.028, -0.019)
23 Term	(-0.004, -0.011, -0.016, -0.015, -0.005, 0.013, 0.039, 0.068, 0.097, 0.122, 0.138, 0.148, 0.138, 0.122, 0.097, 0.068, 0.039, 0.013, -0.005, -0.015, -0.016, -0.011, -0.004)

The corresponding gain functions, which represent the strength remaining at various cycle lengths after the filter is applied, are below:

Figure 4: 5 Term Symmetric Henderson Trend Filter	Figure 5: 7 Term Symmetric Henderson Trend Filter

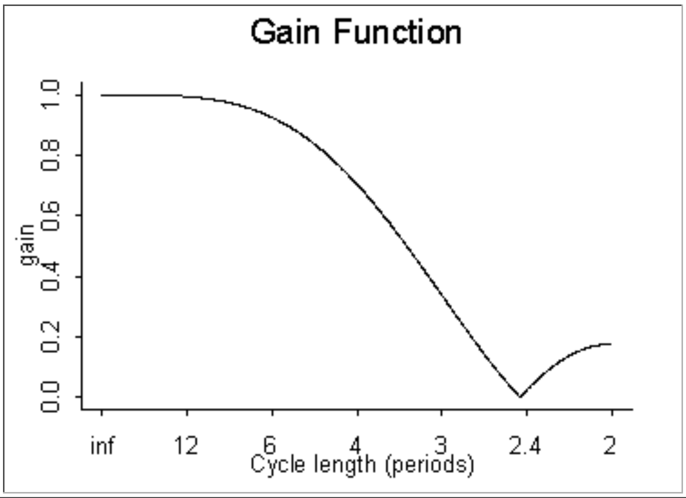

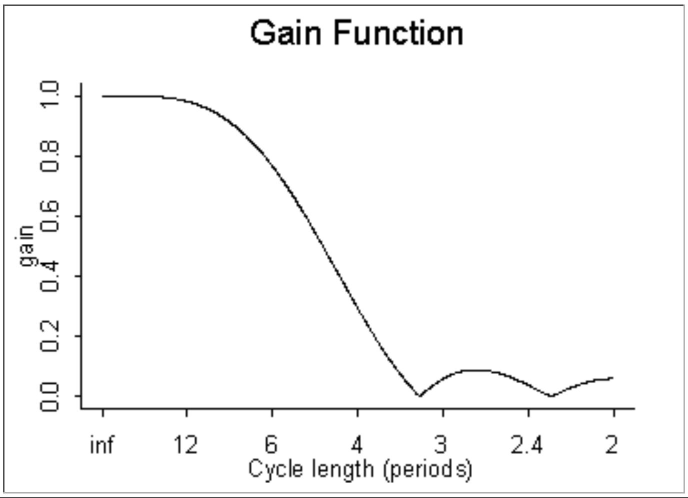

Figure 6: 9 Term Symmetric Henderson Trend Filter

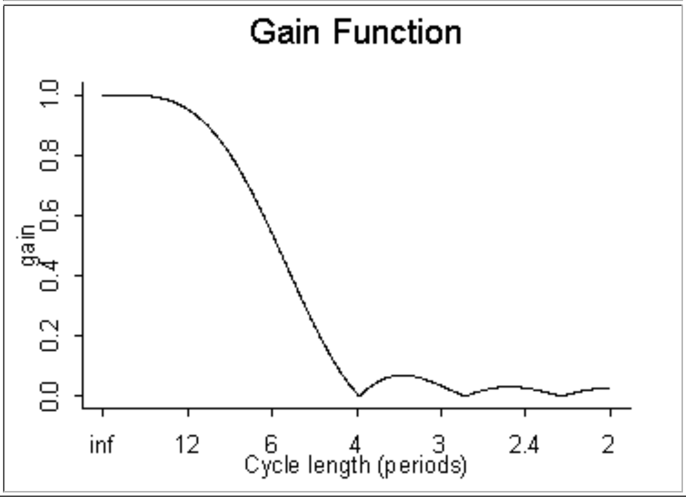

Figure 7: 13 Term Symmetric Henderson Trend Filter

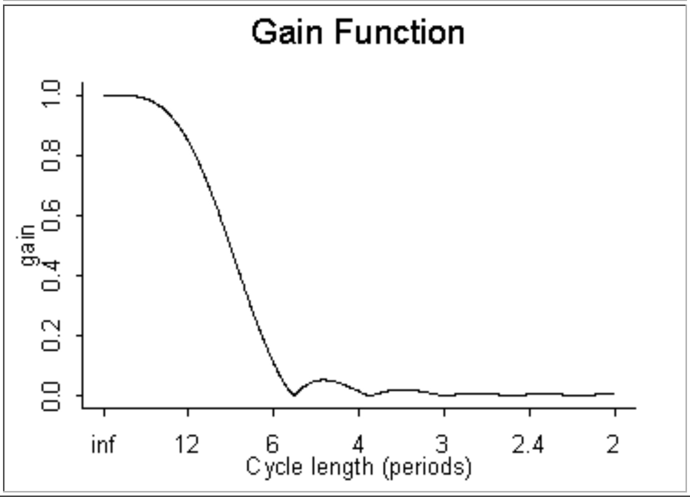

Figure 8: 23 Term Symmetric Henderson Trend Filter

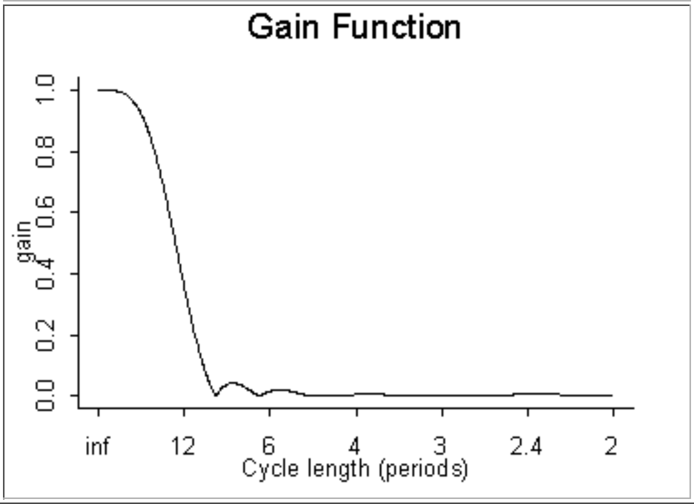

In general, the longer the trend filter, the smoother the resulting trend, as is evident from a comparison of the gain functions above. A 5 term Henderson reduces cycles of about 2.4 periods or less by at least 80%, while a 23 term Henderson reduces cycles of about 8 periods or less by at least 90%. In fact a 23 term Henderson filter completely removes cycles of less than 4 periods.

Henderson moving averages also dampen the seasonal cycles to varying degrees. However the gain functions in Figures 4-8 show that annual cycles in monthly and quarterly series are not dampened significantly enough to justify applying a Henderson filter directly to original estimates. This is why they are only applied to a seasonally adjusted series, where the calendar related effects have already been removed with specifically designed filters.

Figure 9 shows the smoothing effects of applying a Henderson filter to a series:

Figure 9: 23-Term Henderson Filter - Value of Non-residential Building Approvals

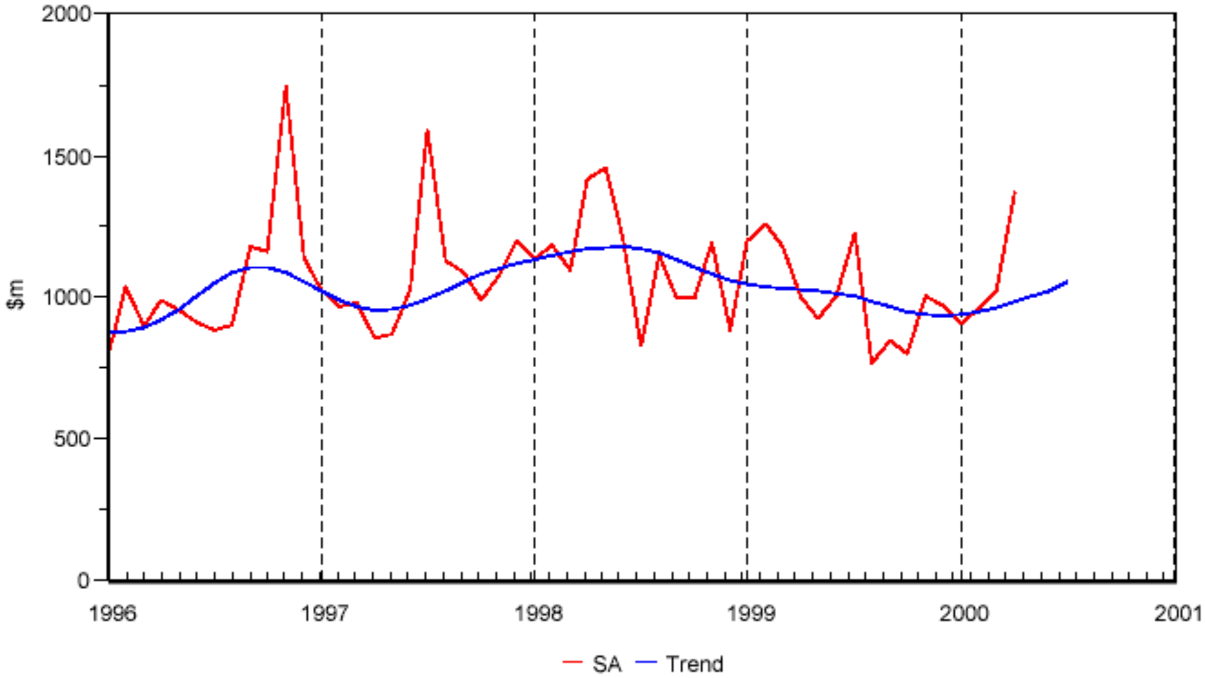

HOW DO WE DEAL WITH THE END POINT PROBLEM?

The symmetric Henderson filter can only be applied to regions of data that are sufficiently far away from the ends of the series. For example the standard 13 term Henderson can only be applied to monthly data that is at least 6 observations from the start or end of the data. This is because the filter smoothness the series by taking a weighted average of the 6 terms on either side of the data point as well as the point itself. If we attempt to apply it to a point that is less than 6 observations from the end of the data, then there is not enough data available on one side of the point to calculate the average.

To provide trend estimates of these data points, a modified or asymmetric moving average is used. Calculation of asymmetric Henderson filters can be generated by a number of different methods which produce similar, but not identical results. The four main methods are the Musgrave method, the Minimisation of the Mean Square Revision method, the Best Linear Unbiased Estimates (BLUE) method, and the Kenny and Durbin method. Shiskin et. al (1967) derived the original asymmetric weights for the Henderson moving average which are used within the X11 packages.

Consider a time series where the last observed data point occurs at time N. Then a 13 term symmetric Henderson filter cannot be applied to data points which are measured at any time after and including time N-5. For all these points, an asymmetric set of weights must be used. The following table gives the asymmetric weighting pattern for a standard 13 term Henderson moving average.

Weights							**Data**						

for period	N-12	N-11	N-10	N-9	N-8	N-7	N-6	N-5	N-4	N-3	N-2	N-1	N
N	0	0	0	0	0	0	-0.092	-0.058	0.012	0.12	0.244	0.353	0.421
N-1	0	0	0	0	0	-0.043	-0.038	0.002	0.08	0.174	0.254	0.292	0.279
N-2	0	0	0	0	-0.016	-0.025	0.003	0.068	0.149	0.216	0.241	0.216	0.148
N-3	0	0	0	-0.009	-0.022	0.004	0.066	0.145	0.208	0.23	0.201	0.131	0.046
N-4	0	0	-0.011	-0.022	0.003	0.067	0.145	0.21	0.235	0.205	0.136	0.05	-0.018
N-5	0	-0.017	-0.025	0.001	0.066	0.147	0.213	0.238	0.212	0.144	0.061	-0.006	-0.034
N-6	-0.019	-0.028	0	0.066	0.147	0.214	0.24	0.214	0.147	0.066	0	-0.028	-0.019

The asymmetric 13 term Henderson filters do not remove or dampen the same cycles as the symmetric 13 term Henderson filter. In fact the asymmetric weighting pattern used to estimate the trend at the last observation amplifies the strength of 12 period cycles. Also asymmetric filters produce some time phase shifting.

WHAT ARE SEASONAL MOVING AVERAGES?

Almost all of the data investigated by the ABS have seasonal characteristics. Since the Henderson moving averages used to estimate the trend series do not eliminate seasonality, the data must be seasonally adjusted first using seasonal filters.

A seasonal filter has weights which are applied to same period over time. An example of the weighting pattern for a seasonal filter would be:

(1/3, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1/3, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1/3)

where, for instance, a weight of one third is applied to three consecutive Januarys.

Within X11, a range of seasonal filters are available to choose from. These are a weighted 3-term moving average (ma) S_{3x1}, weighted 5-term ma S_{3x3}, weighted 7-term ma S_{3x5}, and a weighted 11-term ma S_{3x9}.

The weighting structure of weighted moving averages of the form, S_{nxm}, is that a simple average of m terms calculated, and then a moving average of n of these averages is determined. This means that n+m-1 terms are used to calculate each final smoothed value.

For example, to calculate an 11-term S_{3x9}, a weight of 1/9 is applied to the same period in 9 consecutive years. Then a simple 3 term moving average is applied across the averaged values:

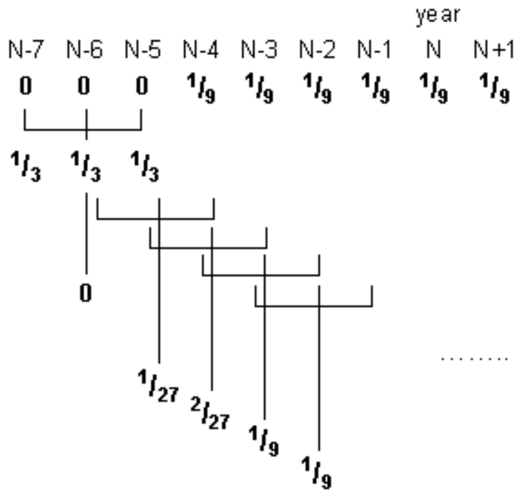

This gives a final weighting pattern of (1/27, 2/27, 1/9, 1/9, 1/9, 1/9, 1/9, 1/9, 1/9, 2/27, 1/27).

The gain function for an 11 term seasonal filter, S_{3x9}, looks like:

Figure 10: Gain Function for 11 Term (S_{3x9}) Seasonal Filter

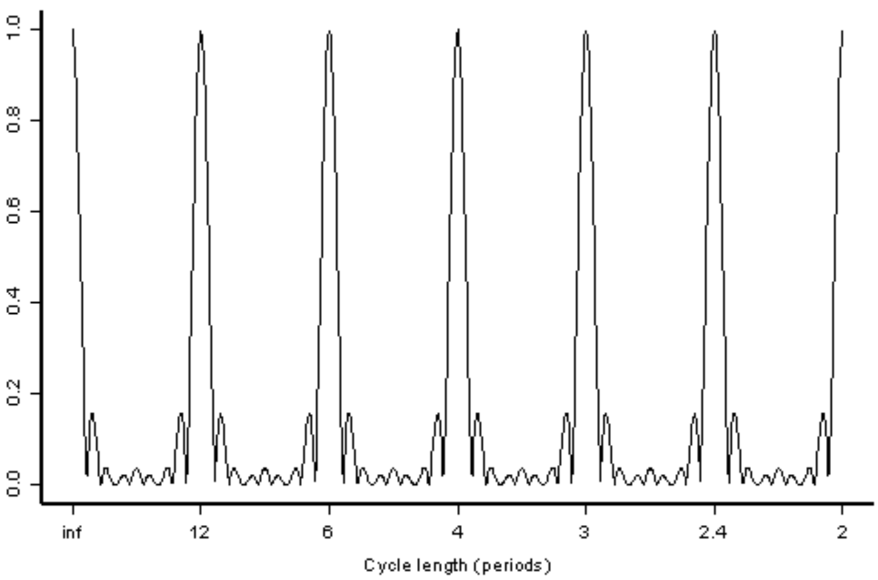

Applying a seasonal filter to data will generate an estimate of the seasonal component of the time series, as it preserves the strength of seasonal harmonics and dampens cycles of non-seasonal lengths.

Asymmetric seasonal filters are used at the ends of the series.

WHY ARE TREND ESTIMATES REVISED?

At the current end of a time series, it is not possible to use symmetric filters to estimate the trend because of the end point problem. Instead, asymmetric filters are used to produce provisional trend estimates. However, as more data becomes available, it is possible to recalculate the trend using symmetric filters and improve the initial estimates. This is

known as a trend revision.

HOW MUCH DATA IS REQUIRED TO OBTAIN ACCEPTABLE SEASONALLY ADJUSTED ESTIMATES?

If a time series exhibits relatively stable seasonality and is not dominated by the irregular component, then 5 years of data can be considered an acceptable length to derive seasonally adjusted estimates from. For a series that shows particularly strong and stable seasonality, a crude adjustment can be made with 3 years of data. It is generally preferable to have at least 7 years of data for a normal time series, to precisely identify seasonal patterns, trading day and moving holiday effects, trend and seasonal breaks, as well as outliers.

HOW DO THE TWO SEASONAL ADJUSTMENT PHILOSOPHIES COMPARE?

Model based approaches allow for the stochastic properties (randomness) of the series under analysis, in the sense that they tailor the filter weights based on the nature of the series. The model's capability for accurately describing the behaviour of the series can be evaluated, and statistical inferences for the estimates are available based on the assumption that the irregular component is white noise.

Filter based methods are less dependent on the stochastic properties of the time series. It is the time series analyst's responsibility to select the most appropriate filter from a limited collection for a particular series. It is not possible to perform rigorous checks on the adequacy of the implied model and exact measures of precision and statistical inference are not available. Therefore, a confidence interval cannot be built around the estimate.

The following diagrams compare the presence of each of the model components at the seasonal frequencies for the two seasonal adjustment philosophies. The x axis is the period length of the cycle and the y axis represents the strength of the cycles which comprise each component:

Figure 11: Comparison of the two seasonal adjustment philosophies

Filter Based

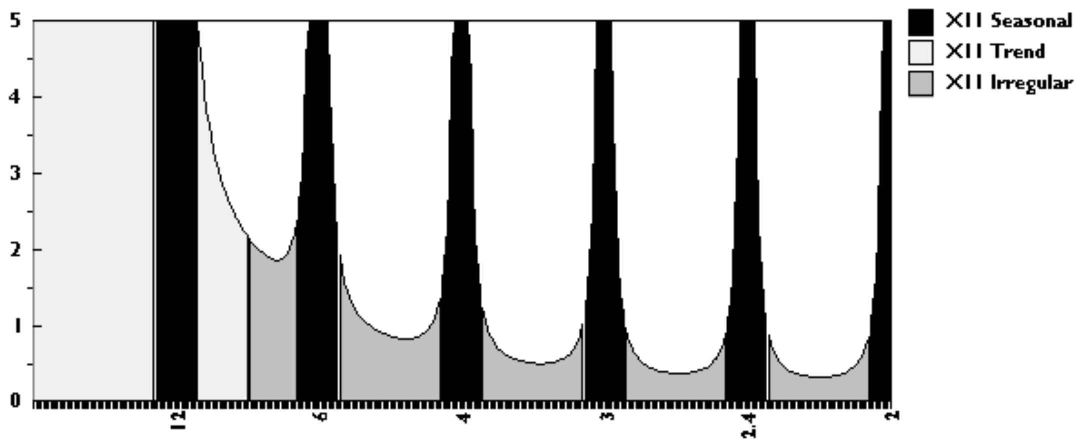

Model Based

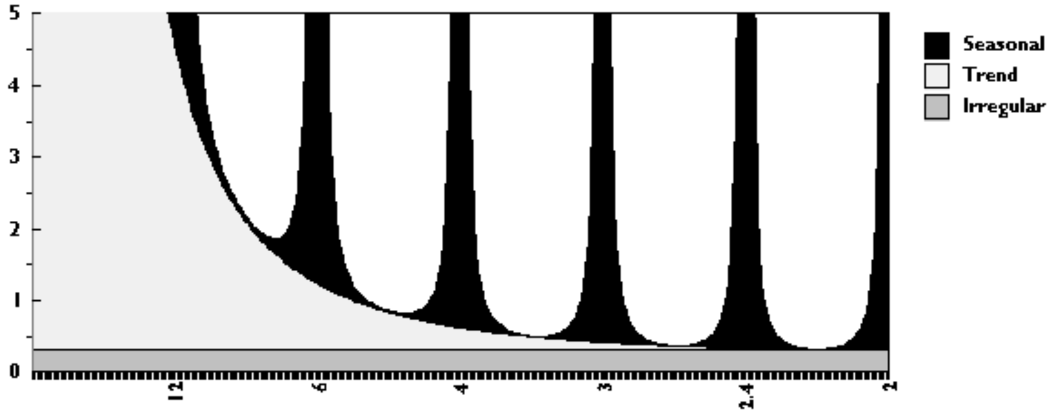

Filter based methods assume that the each component exists only a certain cycle lengths. The longer cycles form the trend, the seasonal component is present at seasonal frequencies and the irregular component is defined as cycles of any other length.

Under a model based philosophy, the trend, seasonal and irregular component are present at all cycle lengths. The irregular component is of constant strength, the seasonal component peaks at seasonal frequencies and the trend component is strongest in the longer cycles.

This page first published 14 November 2005, last updated 1 March 2023

Time Series Analysis: Issues With Seasonal Adjustment

MOVING SEASONALITY

Moving seasonality is a form of seasonality that accounts for the variability in the seasonal component of a time series from year to year. As an example, consider Retail sales in the month of January. Post Christmas specials in January have become more and more popular in recent years. This has resulted in a steady increase in sales for the month of January over the years and is reflected by the slowly evolving nature of the seasonal pattern. This is refered to as moving seasonality.

MOVING HOLIDAYS

Moving holidays are holidays which occur each year, but where the exact timing shifts under the Gregorian calendar system. Examples of moving holidays include Easter and Chinese New Year. Easter generally falls in April but can also fall in late March. Its timing affects series such as Tourism because people often holiday in this period. Chinese New Year mostly occurs in February but can also occur in January. **Overseas Arrivals and Departures** series from some Asian countries are affected by this holiday.

TRADING DAY

The trading day effect is related to months having different numbers of each day of the week from year to year. For example, there may be more garage sales in a January with five weekends rather than four.

In each month, there are four weeks and usually an additional one, two or three days. This means that there are always at least 4 Mondays, 4 Tuesdays, 4 Wednesdays etc., but some days will occur 5 times. A 31 day month for instance, comprises four weeks (28 days) plus three extra days. The number and composition of these extra days will affect the data for the month. For example, assume that a certain store is only open during the week. Then for a 30 day month, within which the two additional days are weekdays, the level of activity is likely to be greater than if the two additional days are Saturday and Sunday.

The Trading day effect is estimated by assigning to each day of the week a weight which reflects the level of activity of that day related to others. For a population with equal levels of activity on the five weekdays and no activity on weekends, weights of 1.4 are assigned to each of the weekdays and 0.0 to each of the weekend days. If a multiplicative (additive) model is used, the daily weights sum to 7 (0) and a weight of 1.0 (0) indicates neutral activity level for the particular day.

A regression framework can be used to estimate the trading day effect from the original estimates. i.e. estimate trading day daily weights for different days. If each daily weight is assumed to be a constant (or static) over time, the data from the entire span of the time series can be used to estimate the daily weights for different days of a week. However, a constant daily weight may not reflect the true behaviour of some time series. Consumer behaviour and trading regulations change over time which means that the trading patterns can change over time. A changing trading pattern related to the day of a week is called moving trading effect.

Moving trading day is used to describe a trading day pattern that is evolving or changing slowly through time. Some changes in the trading day patterns of component series are sudden, while others are relatively gradual. Under Moving Trading Day, the trading day weights are not fixed for the entire length of the series. An example of moving trading day is evident in Australian Retail series. Sunday trading has slowly been introduced over the last ten years or so, in most states of Australia. This has resulted in a moving trading day effect in these series.

To capture the moving trading day effect, a rolling window of a segment of a time series is used to estimate the trading day effect over time.
Figure 1 shows the estimated static and moving daily weights for Australian total retail turnover series. The benchmark estimate is the best estimate of the daily weight that can be made given data available for use at a specified point in time. The graph presents the benchmark estimate (blue line), the mean of the benchmark estimate (black line) and the initial estimate (red dot) of the daily weight factors for each

month. This method of presentation allows the user to see how the benchmark daily weight factors have been evolving over time. It can be observed that the daily weight factors for Sunday have shown a steady rise in later years reflecting the gradual introduction of Sunday Trading in Australia.

Figure 1: Daily weights for Australia total retail turnover

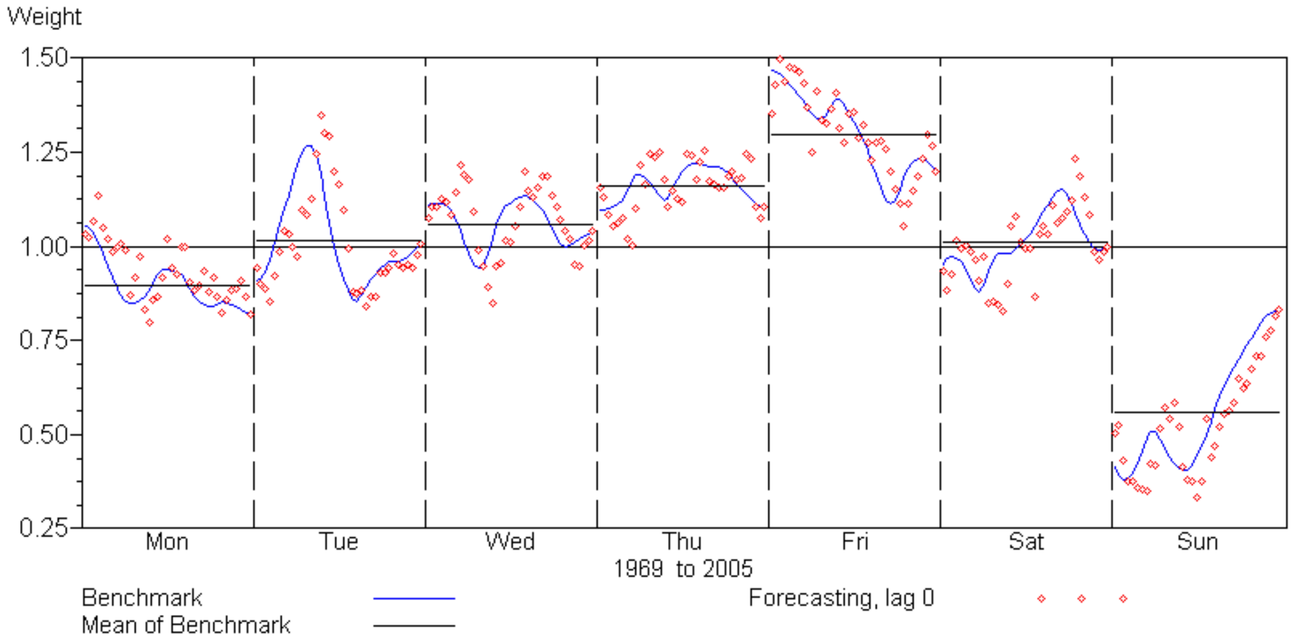

A time series will not exhibit a trading day effect if levels of activity are constant over each day of the week. However, different months have different lengths (28,29,30 and 31 days), hence monthly activity can vary purely because certain months are longer than others. This is known as the length of month effect. If a series has trading day corrections, then these adjustments will include the effect. If there is no trading day effect in a time series, then the length of month effect is accounted for in the seasonal component.

Stock series should not experience a trading day effect since they only measure the level of activity at a certain point in time and are therefore not affected by how many trading days there are in a given period of time.

Why are trading day adjustments rarely made to quarterly series?

Quarters can have 90, 91, or 92 days. Those with 91 days do not experience the trading day effect as this is a multiple of seven, hence they always contain the same number of each day. The others only lose or gain a day, and unless the activity on this day is significant, the effect is nearly impossible to quantify accurately. Trading day effects are rarely seen in ABS quarterly series.

EXTREMES or OUTLIERS

Extremes or outliers are values in a time series that are unusually large or small relative to the other data. They can distort the appearance of the underlying movement of the time series by altering the trend. For this reason, and to improve estimation of the three series components (trend, seasonal and irregular), it is necessary to detect and correct outliers.

For example, consider the Figure 2. The peak in the seasonally adjusted series in late 1994, corresponds to a significantly large irregular value. Because it has not been corrected, it is quite obviously distorting the trend around that time point.

Figure 2: Quarterly Imports of Industrial Transport Equipment

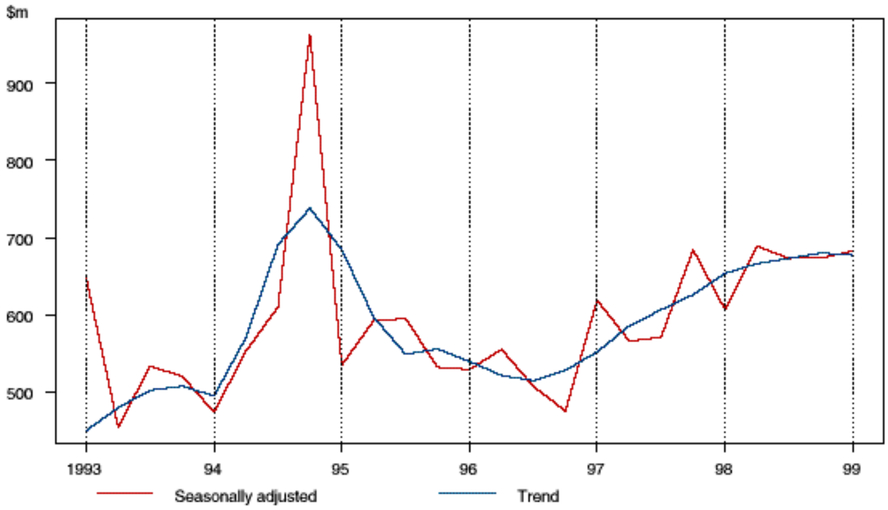

TREND AND SEASONAL BREAKS

Although a time series is a collection of consistently and rigorously defined data items, it is likely that a series will undergo structural breaks during its span. These can be the result of a change in the population they are measuring or in the way that the population is being measured.

Trend breaks

An abrupt but sustained change in the level of a time series is known as a trend break. This is reflected in at least 6 months or 3 quarters of raised or lowered levels. If the span of anomalous values is shorter than this, they are classified as extreme values.

A trend break may be caused by:

- economic policy decisions
  - tariff reductions
- changes in population behaviour
  - St George building society becoming a bank led to a fall in the number of housing loans by non-bank financial institutions.
- changes in the way a population is measured
  - National Accounts no longer including 'sickness benefits' in the measurement of 'other benefits'

.

Seasonal breaks

Seasonal breaks are abrupt changes in the seasonality of a series, which do not affect the level of the series. They may be caused by changes in coverage of the survey, social traditions, administrative practices or technological innovations.

REVISIONS TO TIME SERIES

Estimates of time series components may be revised over time. This can occur for the original, seasonally adjusted and trend estimates. For example, the trend estimate for unemployment for a particular month will change from month to month as new original estimates become available. The revisions to seasonally adjusted and trend estimates will occur until there is enough original estimates available to use symmetric filters to calculate the seasonal and trend components for that month. Typically, such revisions will reduce over time and become negligible after a few months. However, this will depend on the nature of the time series.

FORWARD FACTORS VERSUS CONCURRENT ANALYSIS

There are two approaches to deriving seasonal and trading day factors:

Forward factors rely on an annual analysis of the latest available data to determine the seasonal and trading day factors that will applied to the data received during the forthcoming year.

Concurrent analysis involves re-estimating seasonal factors as each new data point becomes available. Obviously this method is more computationally intense than the forward factor method, but the seasonal factors will be more responsive to dynamic changes.

Currently, the ABS uses forward factor analysis for most time series. However, it has introduced concurrent adjustment into some series and is intending to implement this approach with the majority of ABS time series.

DIRECT (DISAGGREGATE) VERSUS INDIRECT (AGGREGATE) METHODS OF ADJUSTMENT

Sometimes we may deal with series which are related in an aggregative way. For example, we may have data relating to some activity for each individual state, but we would like to obtain a seasonally adjusted series for the Australian total. There are two ways in which we can do this.

Under an indirect (aggregate) method of adjustment, we seasonally adjust each of the lower component series individually, then sum all the values to obtain the seasonally adjusted series for the total.

Directly (disaggregatively) adjusting a series involves summing all the original series to form a total series and then seasonally adjusting the total series directly.

How do we decide which method of adjustment to use?

If the component series each have very different seasonal patterns, then indirect seasonal adjustment is preferable. However if seasonality is minor and difficult to identify in the individual series, then using direct seasonal adjustment may remove any residual seasonality from the aggregate series.

This page first published 14 November 2005, last updated 1 March 2023

Time Series Analysis: Seasonal Adjustment Methods

HOW DO X11 STYLE METHODS WORK?

Filter based methods of seasonal adjustment are often known as X11 style methods. These are based on the 'ratio to moving average' procedure described in 1931 by Fredrick R. Macaulay, of the National Bureau of Economic Research in the US. The procedure consists of the following steps:

1) Estimate the trend by a moving average
2) Remove the trend leaving the seasonal and irregular components
3) Estimate the seasonal component using moving averages to smooth out the irregulars.

Seasonality generally cannot be identified until the trend is known, however a good estimate of the trend cannot be made until the series has been seasonally adjusted. Therefore X11 uses an iterative approach to estimate the components of a time series. As a default, it assumes a multiplicative model.

To illustrate the basic steps involved in X11, consider the decomposition of a monthly time series under a multiplicative model.

Step 1: Initial estimate of the trend

A symmetric 13 term (2x12) moving average is applied to an original monthly time series, O_t, to produce an initial estimate of the trend T_t. The trend is then removed from the original series, to give an estimate of the seasonal and irregular components.

$$\frac{O_t}{\hat{T}_t} = \frac{T_t \times S_t \times I_t}{\hat{T}_t} \approx S_t \times I_t$$

Six values at each end of the series are lost as a result of the end point problem - only symmetric filters are used.

Step 2: Preliminary estimate of the seasonal component

A preliminary estimate of the seasonal component can then be found by applying a weighted 5 term moving average (S_{3x3}) to the $S_t.I_t$ series for each month separately. Although this filter is the default within X11, the ABS uses 7 term moving averages (S_{3x5}) instead. The seasonal components are adjusted to add to 12 approximately over a 12 month period, so that they average to 1 in order to ensure that the seasonal component does not change the level of the series (does not affect the trend). The missing values at the ends of the seasonal component are

replaced by repeating the value from the previous year.

Step 3: Preliminary estimate of the adjusted data

An approximation of the seasonally adjusted series is found by dividing the estimate of the seasonal from the previous step into the original series:

$$\frac{O_t}{\hat{S}_t} = \frac{T_t \times S_t \times I_t}{\hat{S}_t} \approx T_t \times I_t$$

Step 4: A better estimate of the trend

A 9, 13 or 23 term Henderson moving average is applied to the seasonally adjusted values, depending on the volatility of the series (a more volatile series requires a longer moving average), to produce an improved estimate of the trend. The resulting trend series is divided into the original series to give a second estimate of the seasonal and irregular components.

$$\frac{O_t}{\hat{T}_t} = \frac{T_t \times S_t \times I_t}{\hat{T}_t} \approx S_t \times I_t$$

Asymmetric filters are used at the ends of the series, hence there are no missing values like in step 1.

Step 5: Final estimate of the seasonal component

Step two is repeated to obtain a final estimate of the seasonal component.

Step 6: Final estimate of the adjusted data

A final seasonally adjusted series is found by dividing the second estimate of the seasonal from the previous step into the original series:

$$\frac{O_t}{\hat{S}_t} = \frac{T_t \times S_t \times I_t}{\hat{S}_t} \approx T_t \times I_t$$

Step 7: Final estimate of the trend

A 9, 13 or 23 term Henderson moving average is applied to the final estimate of the seasonally adjusted series, which has been corrected for extreme values. This gives an improved and final estimate of the trend. In more advanced versions of X11 (such as X12ARIMA and SEASABS), any odd length Henderson moving average can be used.

Step 8: Final estimate of the irregular component

The irregulars can then be estimated by dividing the trend estimates into the

seasonally adjusted data.

$$\frac{T_t \times I_t}{\hat{T}_t} \approx I_t$$

Obviously these steps will depend on which model (multiplicative, additive and pseudo-additive) is chosen within X11. There are also small differences in the steps in X11 between various versions.

An additional step in estimating the seasonal factors, is to improve the robustness of the averaging process, by modification of the SI values for extremes. For more information on the major steps involved, refer to section 7.2 of the **Information paper: An Introductory Course on Time Series Analysis - Electronic Delivery**.

WHAT ARE SOME PACKAGES USED TO PERFORM SEASONAL ADJUSTMENT?

X11

The most commonly used seasonal adjustment packages are those in the X11 family. X11 was developed by the U.S Bureau of Census and began operation in the United States in 1965. It was soon adopted by many statistical agencies around the world, including the ABS. It has been integrated into a number of commercially available software packages such as SAS and STATISTICA. It uses filters to seasonally adjust data and estimate the components of a time series.

X11ARIMA

The X11 method involves applying symmetric moving averages to a time series in order to estimate the trend, seasonal and irregular components. However at the end of the series, there is insufficient data available to use symmetric weights – the ‘end-point’ problem. Consequently, either asymmetric weights are be used, or the series must be extrapolated.

The X11ARIMA method, developed by Statistics Canada in 1980 and updated in 1988 to X11ARIMA88, uses Box Jenkins AutoRegressive Integrated Moving Average (ARIMA) models to extend a time series. Essentially, the use of ARIMA modelling on the original series helps reduce revisions in the seasonally adjusted series so that the effect of the end-point problem is reduced.

X11ARIMA88 also differs from the original X11 method in its treatment of extreme values. It can be obtained by contacting Statistics Canada.

X12ARIMA

In the late 1990's,the U.S. Census Bureau released X12ARIMA. It uses regARIMA models (regression models with ARIMA errors) to allow the user to extend the series with forecasts and preadjust the series for outlier and calendar effects before seasonal adjustment takes place. X12ARIMA can be obtained from the Bureau.

SEATS/TRAMO

Developed by Victor Gomez and Augustín Maravall, SEATS (Signal Extraction in ARIMA Time Series) is a program which estimates and forecasts the trend, seasonal and irregular components of a time series using signal extraction techniques applied to ARIMA models. TRAMO (Time Series Regression with ARIMA Noise, Missing Observations and Outliers) is a companion program for estimation and forecasting of regression models with ARIMA errors and missing values. It is used to preadjust a series, which will then be seasonally adjusted by SEATS. To freely download the two programs from the internet, contact the Bank of Spain: www.bde.es/homee.htm

DEMETRA

Eurostat has focuses on two seasonal adjustment methods: TRAMO/SEATS and X12ARIMA. Versions of these programs have been implemented in a single interface, called "DEMETRA". This facilitates the application of these techniques to large scale sets of time series. DEMETRA contains two main modules: seasonal adjustment and trend estimation with an automated procedure (e.g. for inexperienced users or for large-scale sets of time series), and with a user-friendly procedure for detailed analysis of single time series.

WHAT ARE THE TECHNIQUES EMPLOYED BY THE ABS TO DEAL WITH SEASONAL ADJUSTMENT?

The main tool used in the Australian Bureau of Statistics is SEASABS (SEASonal analysis, ABS standards). SEASABS is a seasonal adjustment software package with a core processing system based on X11 and X12ARIMA. SEASABS is a knowledge based system which can aid time series analysts in making appropriate and correct judgements in the analysis of a time series. SEASABS is one part of the ABS seasonal adjustment system. Other components include the ABSDB (ABS information warehouse) and FAME (Forecasting, Analysis and Modelling Environment, used to store and manipulate time series data).

HOW DOES SEASABS WORK?

SEASABS performs four major functions:

- Data review
- Seasonal reanalysis of time series
- Investigation of time series
- Maintenance of time series knowledge

SEASABS allows both expert and client use of the X11 method (which has been enhanced significantly by the ABS). This means that a user does not need detailed knowledge of the X11 package to appropriately seasonally adjust a time series. An intelligent interface guides users through the seasonal analysis process, making suitable choices of parameters and adjustment methods with little or no guidance necessary on the users part.

The basic iteration process involved in SEASABS is:

1) Test for and correct seasonal breaks.
2) Test for and remove large spikes in the data.
3) Test for and correct trend breaks.
4) Test for and correct extreme values for seasonal adjustment purposes.
5) Estimate any trading day effect present.
6) Insert or change moving holiday corrections.
7) Check moving averages (trend moving averages, and then seasonal moving averages).
8) Run X11.
9) Finalise the adjustment.

SEASABS keeps records of the previous analysis of a series so it can compare X11 diagnostics over time and 'knows' what parameters led to the acceptable adjustment at the last analysis. It identifies and corrects trend and seasonal breaks as well as extreme values, inserts trading day factors if necessary, and allows for moving holiday corrections.

HOW DO OTHER STATISTICAL AGENCIES DEAL WITH SEASONAL ADJUSTMENT?

- Statistics New Zealand uses X-13ARIMA-SEATS
- Office of National Statistics, UK uses X-12-ARIMA
- Statistics Canada is phasing out X-11ARIMA (as of November 2015) and replacing it with X-12-ARIMA
- U.S Census Bureau uses X-13ARIMA-SEATS
- Eurostat uses Demetra

This page first published 14 November 2005, last updated 1 March 2023

Time Series Analysis: Further Reading

If you would like to learn more about Time Series Analysis, here are some suggestions for further reading material.

Australian Bureau of Statistics (1999a). "The new method for seasonally adjusting crop production data" in Australian Economic Indicators, July 1999 (cat. no. 1350.0). **Australian Bureau of Statistics**, Canberra, Australia.

Australian Bureau of Statistics (1999b). Introduction of Concurrent Seasonal Adjustment into the Retail Trade Series (cat. no 8514.0). **Australian Bureau of Statistics**, Canberra, Australia.

Australian Bureau of Statistics (2003). Information Paper: A Guide to Interpreting Time Series - Monitoring Trends (cat. no. 1349.0). **Australian Bureau of Statistics**, Canberra, Australia.

Bell, P. (1999). The Impact of Sample Rotation Patterns and Composite Estimation on Survey Outcomes, Working Papers in Econometrics and Applied Statistics, No. 99/1, (cat. no. 1351.0). **Australian Bureau of Statistics**, Canberra, Australia.

Bell, W.R. and Hillmer, S.C. (1984). Issues Involved with the Seasonal Adjustment of Economic Time Series, **Journal of Business and Economic Statistics**, 2, 291-320.

Binder, D.A. and Hidiroglou M. (1988). Sampling in Time, in Handbook of Statistics, Vol. 6, ed. By P.R. Krishaiah and C.R. Rao, Elsevion Science Publishers, B.V., 187-211.

Box, G.E.P; Jenkins, G.M. (1970). Time Series Analysis: Forecasting and Control, published by Holden-Day.

Cannon, J. (2000). Diagnostic Measures for Comparing Direct and Aggregative Seasonal Adjustments, Working Papers in Econometrics and Applied Statistics, No. 2000/1, (cat. no. 1351.0). **Australian Bureau of Statistics**, Canberra, Australia.

Cannon, J. and Van Halderen, G. (2000). Aggregation and ABS Time Series. Methodology Advisory Committee Research Paper, July 1999 (cat. no. 1352.0.55.030). **Australian Bureau of Statistics**, Canberra, Australia.

Cantwell, P.J., and Caldwell, C.V. (1998). Examining the Revisions in Monthly Retail and Wholesale Trade Surveys Under a Rotating Panel Design. **Journal of Official Statistics**, Vol. 14, No. 1, 47-59.

Chatfield, C. (1996). The analysis of time series : an introduction. (5th edition). London, Chapman and Hall.

Cleveland, R.B., Cleveland, W.S., McRae, J.E., and Terpenning, I. (1990). STL: A

Seasonal-Trend Decomposition Procedure Based on Loess, **Journal of Official Statistics**, Vol. 6, No. 1., 3-73.

Dagum, E.B. (1980). The X11ARIMA Seasonal Adjustment Method. Ottawa: Statistics Canada, cat. 12-564E.

Dagum, E.B., Chhab, N. and Chiu, K. (1996). Derivation and Properties of the X11ARIMA and Census X11 Linear Filters. **Journal of Official Statistics**, 12, No. 4, 329-347.

Doherty, M.(1992). The Surrogate Henderson Filters in X11, **Statistics New Zealand**, Working Paper.

Doherty, M. (2001). The Surrogate Henderson Filters in X11, **Australia & New Zealand Journal of Statistics**, Vol 43, No. 4, 385-392.

Durbin, J. (2000). The Foreman lecture: The state space approach to time series analysis and its potential for official statistics (with Discussion), **Australia & New Zealand Journal of Statistics**, Vol 42, No. 1, 1-24.

Findley, D.F., Monsell, B.C., Bell, W.R., Otto, M.C. and Chen B. (1998). New Capabilities and Methods of the X-12-ARIMA Seasonal-Adjustment Program. **Journal of Business and Economic Statistics**, Vol. 16, No. 2., 127-177.

Gray, A. and Thomson, P. (1996a). Design of moving-average trend filters using fidelity, smoothness and minimum revisions criteria, Bureau of the Census, RR96/01.

Gray, A. and Thomson, P. (1996b). Design of moving-average trend filters using fidelity and smoothness criteria in Vol 2: Time Series Analysis in Memory of E.J. Hannan. ed. P.Robinson and M. Rosenblatt. Springer Lecture Notes in Statistics 115, 205-219.

Gray, A. and Thomson, P. (1996c). On a family of moving-average trend filters for the ends of series. Proceedings of the American Statistical Association, Section on Survey Research Methods, 1996.

Henderson, R. (1916). Note on Graduation by Adjusted Average. Transactions of the American Society of Actuaries, 17, 43-48.

Harvey, A. (1990). Forecasting, structural time series models, and the Kalman filter. Published by Cambridge University Press.

Kenny, P.B., and Durbin, J. (1982). Local Trend Estimation and Seasonal Adjustment of Economic and Social Time Series. **Journal of the Royal Statistical Society**, Series A, 145, 1-41.

Ladiray, D. and Quenneville, B. (2001). Seasonal Adjustment with the X-11 method, New York: Springer Verlan, Lecture notes in statistics, 158.

Laniel, N. (1985). Design criteria for 13 term Henderson end-weights. Technical Report Working paper TSRA-86-011, **Statistics Canada**, Ottawa K1A 0T6.

Leung C., McLaren C.H., Zhang X~~,~~. (1999). Adjusting for an Easter Proximity Effect, Working Papers in Econometrics and Applied Statistics, No. 99/3, (cat. no. 1351.0). **Australian Bureau of Statistics**, Canberra, Australia. (N.B. This paper is located on the same page as Working Papers in Econometrics and Applied Statistics: No 99/2 Seasonal Adjustment - Comparison of Philosophies, Dec 1999. Go to the downloads tab to find both papers.)

Macaulay, F. (1931). The smoothing of time series. **National Bureau of Economic Research**, New York.

McLaren, C.H. (1999). Designing Rotation Patterns and Filters for Trend Estimation in Repeated Surveys Unpublished PhD. Thesis, **University of Wollongong**, NSW, Australia.

Pena, D., Tiao, G.C. and Tsay, R.S. eds (2001). A Course in Time Series Analysis. Wiley Series in Probability and Statistics. Wiley-Interscience, New York.

Pierce, D.A. (1980). Data Revisions with Moving Average Seasonal Adjustment Procedures. **Journal of Econometrics**, 14, 95-114.

Shiskin, J., Young, A. H. And Musgrave, J.C. (1967). The X11 Variant of the Census Method II Seasonal Adjustment Program. Technical Paper 15, Bureau of the Census, U.S. Department of Commerce, Washington, D.C.

Sutcliffe, A. and Lee, G. (1995). Seasonal Analysis and Sample Design. Paper presented at the Conference of Survey Measurement and Process Quality, Bristol 1995.

Sutcliffe, A. (1993). X11 Time Series Decomposition and Sampling Errors. Working Papers in Econometrics and Applied Statistics, No. 93/2. **Australian Bureau of Statistics**, Canberra, Australia.

Sutcliffe, A. (1999). Seasonal adjustment: Comparison of Philosophies. Working Papers in Econometrics and Applied Statistics, No. 99/2, (cat. no. 1351.0). **Australian Bureau of Statistics**, Canberra, Australia.

Von Sanden, N. and Zhang, X.M. (2001). Use of concurrent seasonal adjustment for economic time series: the case for Retail Survey. Methodology Advisory Committee Research Paper, June 2001 (cat. no. 1352.0.55.039). **Australian Bureau of Statistics**, Canberra, Australia.

Wallis, K.F. (1982). Seasonal Adjustment and Revision of Current Data: Linear filters for the X11-method. **Journal of the Royal Statistical Society**, Series A, 145, 74-85.

Wei, W.W.S. (1993). Time Series Analysis: Univariate and Multivariate Methods. Published by Addison-Wesley.

Zhang, X.M. and Apted, L. (2008). Temporal Aggregation and Seasonal Adjustment. Methodology Advisory Committee Research Paper, June 2008 (cat. no. 1352.0.55.095). **Australian Bureau of Statistics**, Canberra, Australia.

Zhang, X., McLaren, C.H., Leung, C.C.S. (2001). An Easter proximity effect:

modelling and adjustment. **Australian & N.Z. Journal of Statistics** Vol. 43, No. 3, 269-280.

Zhang, X.M., Von Sanden, N., Menezes, Z. and McLaren, C. (2006). Some Aspects of Turning Point Detection in Seasonally Adjusted and Trend Estimates. Methodology Advisory Committee Research Paper, June 2006 (cat. no. 1352.0.55.079). **Australian Bureau of Statistics**, Canberra, Australia.

Zhang, X.M. and Sutcliffe, A. (2001). Use of ARIMA Models for Improving the Revisions of X-11 Seasonal Adjustment. Methodology Advisory Committee Research Paper, November 2001 (cat. no. 1352.0.55.042). **Australian Bureau of Statistics**, Canberra, Australia.

This page first published 11 November 2005, last updated 1 March 2023